

Frontiers of Research in Advanced Computations

The Laboratory's newly reorganized Institute for Scientific Computing Research fosters collaborative research in advanced computing techniques. Recent work is laying the foundations for innovative and sometimes startling methods in computational physics, massively parallel processing, computer vision, the modeling of human joints, and a range of other applications.



OUR national laboratories have been widely regarded as the undisputed leaders in computational physics. That capability arose because researchers faced enormous intellectual challenges that required inventive solutions on the best hardware.

Today, missions have evolved, budgets are tighter, and we no longer have exclusive access to the best hardware. Yet, Lawrence Livermore National Laboratory remains a unique place where expertise spans a wide range of disciplines and the

cross-fertilization of ideas abounds. When University of California researchers and resources are added to the equation, the possibilities are genuinely exciting.

In 1996, several computational groups at the Laboratory joined the Institute for Scientific Computing Research (ISCR). This was part of an overall project to realign several of Lawrence Livermore's Institutes and Centers in order to advance the strategic goals of the Laboratory, the Department of Energy, and the University of California and to provide a productive environment for university faculty and students within budgetary constraints.

The ISCR now has three components: the Center for Computational Physics, the computational biomechanics group, and the computer vision group. In addition to its principal research projects, the ISCR's outreach activities include:

- Funding collaborative research at university campuses.
- Sponsoring postdoctoral researchers at Livermore.
- Conducting seminars and workshops.
- Arranging consultant and guest activities.

- Pursuing technology transfer initiatives.
- Conducting work for others.

To accomplish its mission, the ISCR assesses Laboratory needs to see what is missing in existing computational methods, and it often charts new directions. Its activities represent a balance between complementing, but not duplicating, ongoing programs at Lawrence Livermore and exploring entirely new concepts in scientific computing.

Using the considerable resources of the Laboratory and the University, the ISCR is developing techniques at the cutting edge of computation. Its inventions include better fluid- and plasma-transport models, new solid-state low-frequency magnetic field solvers, new concepts in

massively parallel processing, real-time object and motion recognition systems, and improved models of human joint dynamics.

Grid and Particle Hydrodynamics

As part of the ISCR, the Center for Computational Physics (formerly within the Plasma Physics Research Institute) develops simulations that help researchers study the behavior of plasmas, which are highly or completely ionized gases, and other physics phenomena. One new algorithm, Grid and Particle Hydrodynamics (GaPH, pronounced “gaff”), is a computational tool for studying the complex behaviors of a plasma or a gas. GaPH was developed to help scientists and engineers understand

more about the chemistry of systems with complex geometries and to do so at far less cost than that of other methods. We can understand much of its purpose through a simple example of the kind of problem GaPH was designed to address.

Figure 1 shows data about two localized gas “puffs” that expand into one another, colliding and interpenetrating. Figure 1a is a snapshot of two sharply defined spikes of gas that are slightly separated in space. Most of the gas particles are moving slowly, but a few are about to move quickly in one direction or another (those with large negative or positive velocity). It is like a snapshot of two large groups of people in Grand Central Station—some people are standing still, others are strolling, some are rushing to catch a

train, but all are temporarily frozen in time. Then the two gas puffs “splash” into each other. Figure 1b does not show collisions; Figure 1c is a simulation with collisions included. When gas puffs interpenetrate, many steepening pressure waves are formed that can become complex and turbulent. Simulating those turbulent waves is precisely the type of problem that Laboratory scientists need to address in studying z-pinch (structures used to generate x rays) and interpenetrating plasmas in National Ignition Facility targets and in weapons systems.

The current approach at the Laboratory is to use a fluid treatment (hydrodynamics codes) to study these systems. But when these systems are driven hard by external forces, a pure fluid treatment fails to recover important

features. Mixing and turbulence build and grow so fast that collisions do not relax the system back to consistency with fluid models.

A common alternative that overcomes some of these difficulties is the particle-in-cell (PIC) method, which groups many similar particles together into macroparticles and follows their interactions using discrete time steps. Even so, today’s computers can follow only a small fraction of the events of interest. A PIC collection can hint at the essential features, but it takes experienced eyes to see signals in the noise.

GaPH is a better tool for understanding gas and plasma behavior. GaPH can model systems in which gases or plasmas either do not collide (Figure 1b) or do collide (Figure 1c) with one another. The code can be used to reconstruct the distribution of real particles at all points in space at any time of interest.

GaPH starts with a relatively small number of “smart” particles, each of which is a lump of fluid representing many (perhaps billions of) real gas molecules. Each GaPH representation has velocities in all directions based on the internal dynamics within the lump. Over time, individual superparticles expand due to their internal energy (pressure) and velocity.

GaPH is unique in that it continuously allows new superparticles to be “born” so that they will be available where interesting things are happening (Figure 2). Conversely, superparticles with overlapping properties can be merged. By eliminating redundant representations, GaPH wastes less computational effort and focuses more efficiently on the most relevant collisions or events. For example, a one-dimensional GaPH simulation needs only 400 superparticles rather than the 20,000 macroparticles required for a standard PIC problem. The important points are that GaPH allows investigators to spend their computer resources on




Figure 2. Imagine that a simulated superparticle (a) is a sugar cube containing billions of atoms. As the cube melts and spreads out (b), GaPH continuously adds new particles to the simulation (c) to account for what is happening at the edges.

those parts of a problem that require the most scrutiny, and GaPH can account for the interactions that escape standard fluid treatments.

Even though GaPH is a new concept, it already appears to be the best tool to address interpenetration in turbulent systems with low rates of collision. The next step is to extend GaPH to three-dimensional representations and to introduce more realistic physics.

Electromagnetic Modeling

The Center for Computational Physics is also a center of excellence for the computer modeling of low-frequency electromagnetic phenomena (often called Darwin models after their originator) in plasmas and magnetic materials. In these models, high frequencies (light waves) are neglected, thus eliminating considerable computational effort. Such numerical simulations are important for many Laboratory and industrial applications of plasmas.

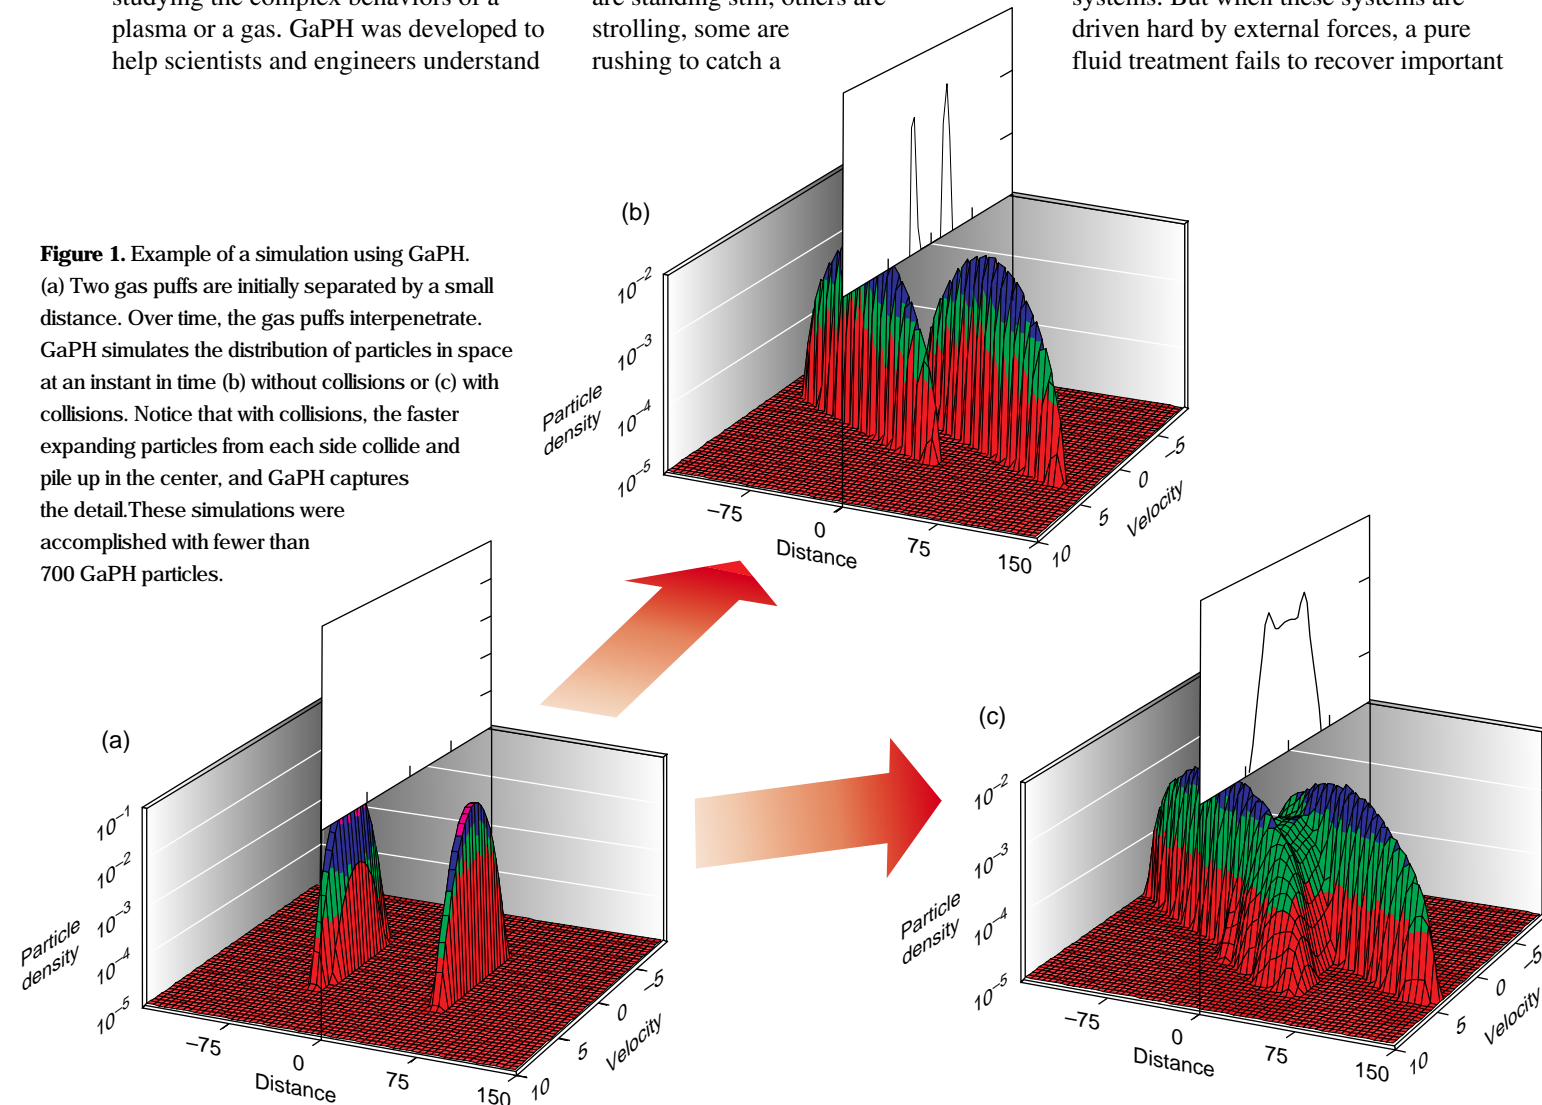
Laboratory programs are concerned with the behavior of plasmas in etching and deposition processes, magnetic fusion, laser fusion, and other areas including defense. Plasmas can be simulated with fluid, PIC, or GaPH techniques, but each one requires a suitable way to calculate the electromagnetic fields that interact with the charged particles of plasma. The ISCR is developing Darwin models that

provide the fields for any of these techniques.¹

In semiconductor wafer etching, an important plasma application, researchers need to design reaction chambers that properly confine plasma, and they need to optimize antennas that can generate a uniform plasma of maximum density across the wafer surface (Figure 3). The ISCR has developed models that describe this process, simulating both resistive heating (similar to current through a wire) and stochastic heating, which depends on the distribution of plasma particles. Institute researchers have concentrated on extending their models to address three-dimensional problems.

Many other applications can benefit from the same types of computational methods. One example is high-speed flywheels that can serve as electromechanical batteries (see the April 1996 issue of *Science & Technology Review*). Another application is in the magnetic recording industry, which has developed a new, giant magneto-resistive material. This material allows changes in resistivity (the detection of which results in a “read”) to be caused by a much smaller change in a magnetic field. Thus, a “bit” can be localized in a much smaller area, enabling information to be packed more densely than on present-day magnetic media. The interactions behind this concept (those

Figure 1. Example of a simulation using GaPH. (a) Two gas puffs are initially separated by a small distance. Over time, the gas puffs interpenetrate. GaPH simulates the distribution of particles in space at an instant in time (b) without collisions or (c) with collisions. Notice that with collisions, the faster expanding particles from each side collide and pile up in the center, and GaPH captures the detail. These simulations were accomplished with fewer than 700 GaPH particles.



between low-frequency magnetic fields and induced electric fields) are a perfect application for the ISCR’s Darwin models.

Massively Parallel Processing

A popular view is that the future of supercomputing will depend on massively parallel computers. The arguments are persuasive, so what is the delay? For one thing, compilers (programs that convert scientific programming language into machine language) do not yet use all the capabilities of the newest hardware. But even when compilers catch up, users still have to reorganize their algorithms and the way they think about solutions to realize the promise of massively parallel computing.




Figure 3. To process wafers as semiconductors, antennas generate plasmas by heating electrons. In this simulation of a two-dimensional plasma-processing chamber, the contours are the magnitude of the electric field driven by the antenna. Yellow represents high electric field intensity; blue indicates low electric field intensity.

Massively parallel processing (MPP) systems can have 100, 1,000, or even more microprocessor-based central processing units (CPUs). During a complex calculation, a problem is broken down into tasks or fragments. The difficulty is that, at some point, the processor assigned to a given task needs information computed elsewhere. In many cases, all parts of the system must talk to every other part before a solution is reached. Slightly stretching the point, it is as if every U.S. citizen had to talk with every other citizen before a candidate was elected President. MPP users worry about data layout across all the processors, synchronization between tasks, data transfer rates, and many other issues.

Three years ago, Institute researchers began exploring MPP techniques to solve linear systems that are the backbone of

many codes used at the Laboratory and other institutions. This work involves collaborations with faculty and students at the University of California campuses at Davis and Los Angeles and with LLNL researchers outside the Institute.

The Institute developed a new linear system to implement alternating-direction-implicit (ADI) methods but found that it also is useful in other areas. ADI codes split a big computational problem into a series of independent linear systems. Rather than giving an entire “line,” or part of the problem, to each processor, each line is split over several processors, and adjacent segments of neighboring lines can now be given to the same processor. Equally important, the domain structure—the spatial partitioning of problem parts to each processor—remains unchanged during the entire solution process.

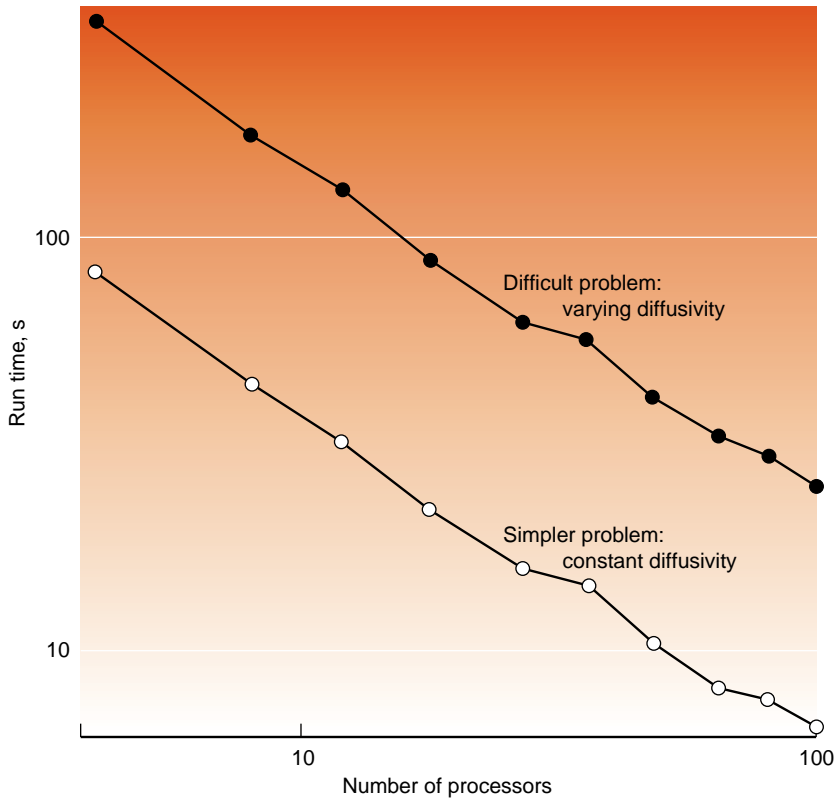


Figure 4. The curves show results for a three-dimensional, dynamic alternating-directional-implicit (ADI) solution of the steady-state diffusion equation involving 864,000 unknowns. Note the large decrease in time as the number of processors increases.

The Institute’s new methods have several advantages (see Figure 4). The part of a problem associated with each processor can be tailored to minimize the information exchange between processors, which is a particularly important issue for PIC or fluid models. The methods make it possible to apply plasma and magnetic-material codes in two and three dimensions with relatively high resolution. It will also be easy to apply the new concepts as MPP multiprocessor technology continues to mature.

Curved-Boundary Modeling

Many computer simulations use irregular mesh elements to represent

structures with curved boundaries. In some cases, mesh points can move with a structure so that the model follows the motion of the structure.

Orthogonal meshes consist of a set of straight lines that intersect at right angles at mesh points. A close look at a curved boundary represented this way reveals a jagged representation (Figure 5a). Although this approach is adequate for certain problems, the computer representation of electromagnetically driven particle motion near such boundaries is incorrect and often unacceptable.

Livermore’s ISCR has developed a new embedded curved-boundary (ECB) method that offers the utility and

flexibility of unstructured meshes while retaining the speed and user-friendly characteristics of orthogonal meshes. Curved boundaries are embedded within an orthogonal mesh, making it possible to model realistic curved boundaries on a computationally convenient mesh. The advantages (Figures 5b through 5d) include much quicker solution of the differential equations required in the vicinity of a curved boundary. As with other boundary models, embedded curved boundaries can also be moved at the user’s discretion to follow the motion of a structure.

This work is closely connected to the Institute’s other efforts in that the ECB method builds more capable

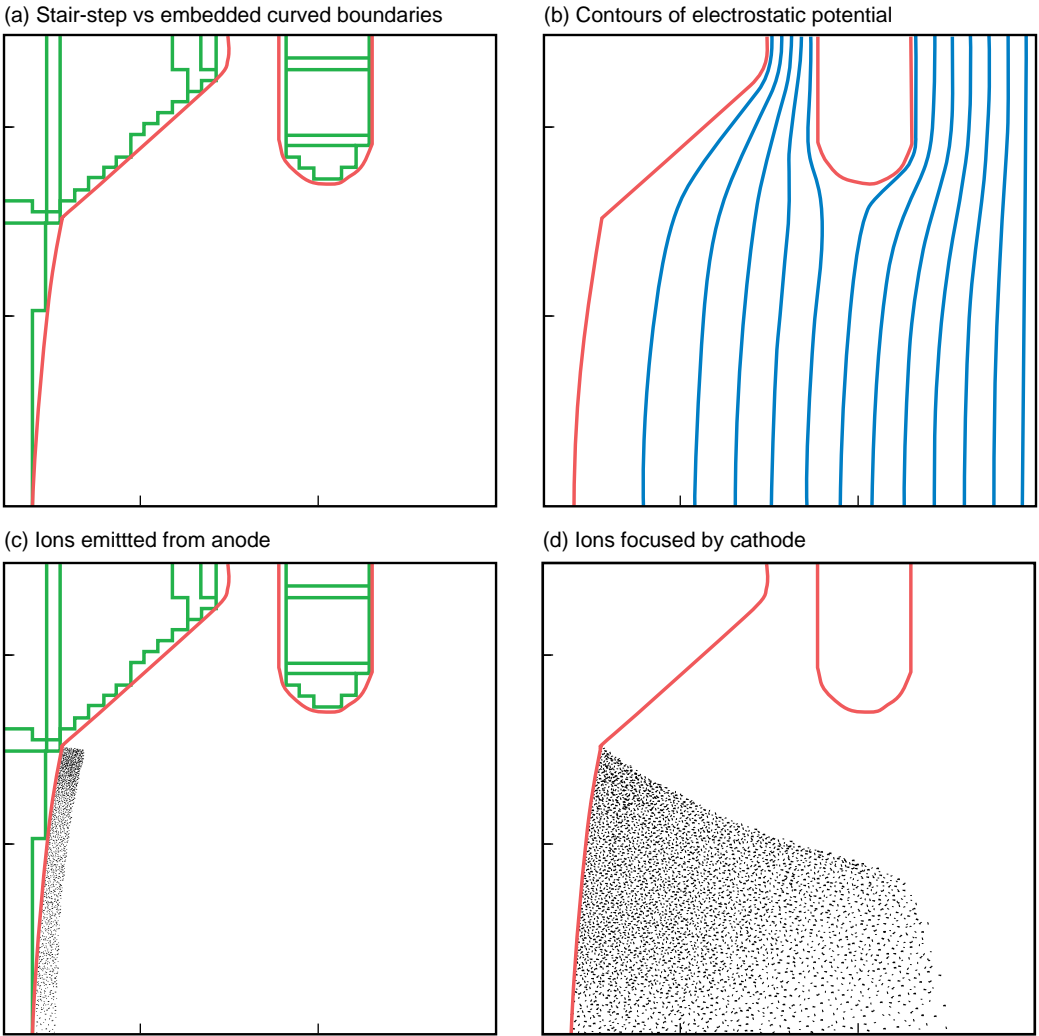


Figure 5. Cross sections of a typical ion injector. (a) Compared to an orthogonal mesh, embedded curved boundaries more accurately represent the actual electrode surfaces. (b) Calculated contours of electrostatic potential using curved boundaries. Ions are (c) emitted from the curved anode and (d) subsequently focused by the extraction cathode.

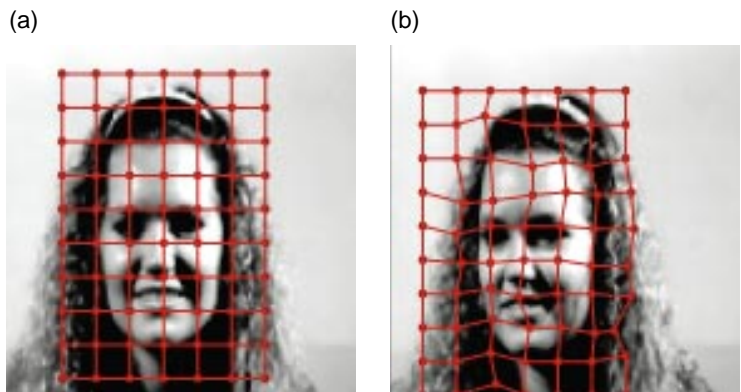


Figure 6. (a) KEN outlines a face to be matched with a grid overlay, here shown as a rectangular grid for clarity, and stores it in its memory. (b) Matching is determined by how closely a new image fits the grids stored in the database.

representations of the differential equations in the vicinity of curved boundaries. The result is that the ECB method can be seamlessly added to the newly improved capability in massively parallel processing. Taken together, GaPH, Darwin models, and ECB methods are adding considerable power to the Laboratory's modeling strength and to the move toward massively parallel implementation.

Computer Vision

Using computers to recognize objects has enormous possibilities in the era of the information superhighway. Automated object and motion recognition can be applied in security and surveillance, medical, defense, and telecommunications applications as well as in a host of other areas. Computerized object recognition would be an invaluable tool for searching image databases. Face recognition could, for example, be used to verify credit cards or other valuable property. An autonomous robot with a recognition system could access places or perform tasks that are impractical for humans.

The ISCR developed a near-real-time face-recognition technology, KEN, which was previously reported on in *Energy & Technology Review*.² As shown in

Figure 6, KEN extracts information about a face in the form of a grid marked with features and stores this model in memory. To recognize a face, KEN compares all face models in its database to the unknown face. After statistical evaluation of similarities and differences, the system rejects poor matches and selects a qualified match if one is found. Using a database of several hundred faces, KEN can identify up to 98.5% of the faces correctly. Industry contacts from TRW and Intel (among others), law enforcement agencies in Europe and California, and the FBI have expressed interest in the technology.

The ISCR's computer vision group is extending KEN to include much larger databases, to organize the databases by comparing stored face models with each other, and to recognize other object classes, such as footprints, signatures, and graphics. They also recently began developing a motion-recognition system featuring a new motion-sensitive silicon retina. This work is a logical extension of KEN, which is based on a comparison of two images. Motion recognition tracks the distortion or changes occurring in a succession of several images.

The ISCR's approach to computer vision incorporates advanced, modular, mix-and-match components in hardware and software. The components are based

on artificial neural networks and neuromorphic engineering concepts, which mimic the structure and activities of the brain. The Institute's work at the forefront of computer-vision research attempts to mimic a type of motion-detection process found in biological visual systems. More specifically, computers can imitate the way specialized neurons in the retina respond to a moving target but do not react when a target is stationary.

ISCR researchers can use either a charge-coupled device (CCD) camera or analog silicon-retina chips as the input sensor (or "eye") for a computer. These chips, developed by a research group at the California Institute of Technology, have improved dynamic range in difficult lighting conditions compared to a CCD camera.

The motion-recognition system being developed will combine a high-resolution silicon retina, a motion-sensitive chip, devices for data capture and processing, and object-oriented software components. **Figure 7** shows some early results from tests of a motion-sensitive chip. A prototype system will be up and running by the fall of 1996.

An important spinoff of motion analysis involves data compression of video sequences. The new method developed by ISCR researchers uses motion-assisted segmentation to yield a higher data-compression factor (up to 350 to 1) of generic video test data with fewer errors than other methods. This method could contribute a component to the MPEG-4 Standard currently under development by the Motion Picture Expert Group (MPEG). MPEG-4 specifically aims at low-bitrate and wireless communication.

Biomechanical Modeling

ISCR researchers in collaboration with the Laboratory's Engineering Directorate are also developing computational models of the structure

and function of human joints. In this work, ISCR researchers begin with very high-resolution scans from individuals and use surface extraction and finite-element techniques to create highly detailed, accurate models of joint dynamics. A three-dimensional, nonlinear, finite-element model (NIKE3D, developed at LLNL for engineering problems) allows the ISCR biomechanics group to address biological problems realistically. Researchers can assess interactions among different types of tissue—including bones, ligaments, tendons, and muscle—when they assign mechanical properties and physiological loads to each structure within a joint.

Biomechanical modeling will lead to a better understanding of repetitive strain injury, degenerative joint diseases, and traumatic injury. A current focus is on applying the joint models to solve problems in the orthopedic industry, specifically to extend the quality and life span of prosthetic joint implants. This biomechanical modeling effort will be the topic of a research highlight in the September issue of *Science & Technology Review*.

ISCR work has also included finding a way to noninvasively monitor blood oxygen in real time and developing Sisal, a functional language that simplifies the programming of parallel supercomputers.

The goal of the Sisal Project is to have the system software automatically manage the machinery and allow the programmer to focus on the problem and its solution. By speeding up the coding process and supporting existing codes written in other languages, ISCR researchers developed a way to make portable parallel computing more practical and affordable than ever before. Two spinoffs of this project are the Massively Parallel Input/Output Project, now in its third year, and the High-Performance Functional Computing Project. In short, Institute researchers do more than simply refine old methods or apply




Figure 7. An almost stationary target (left) yields no signal. Faster motion of a model car (right) traveling from left to right increases the signal seen by a motion-sensitive silicon retina designed at the California Institute of Technology.

them efficiently to new hardware; it is not enough to do a job several times faster on a better machine. Rather, the Institute seeks alternative ways to represent physical information, to bridge the gap between computer science and scientific computing applications, and to reach solutions. When a given project is successful, an entirely new program may be born.

Key Words: advanced computer modeling methods—alternating-direction-implicit (ADI), embedded curved-boundary (ECB), grid and particle hydrodynamics (GaPH), KEN, low-frequency electromagnetic (Darwin); computer vision; Institute for Scientific Computing Research (ISCR); massively parallel processing (MPP).

About the Team



Some of the members of LLNL's **INSTITUTE FOR SCIENTIFIC COMPUTING RESEARCH** are (left to right) Michael A. Lambert, William B. Bateson, Karin Hollerbach, Matthew Gibbons, Dennis W.

Hewett, Louann S. Tung, and Martin Lades. This newly reorganized Institute does collaborative research in advanced computing techniques with programs and clients inside and outside the Laboratory and is currently focusing on innovative computing methods in computational physics, massively parallel processing, computer vision, and biomechanical modeling, among others.

References

1. D. W. Hewett, "Low-Frequency Electromagnetic (Darwin) Applications in Plasma Simulation," *Computational Physics Communications* **84**, 243 (1995); M. R. Gibbons and D. W. Hewett, "The Darwin Direct Implicit Particle-In-Cell (DADIPIC) Method for Simulation of Low Frequency Plasma Phenomena," *Journal of Computational Physics* **128**, 231 (1995).
2. "KEN Project: Real-World Face Recognition," *Energy & Technology Review*, UCRL-52000-94-10 (**October 1994**), pp. 22–23.

For further information contact
Dennis W. Hewett (510) 422-5432
(dhewett@llnl.gov).